

CLAIMS

What is claimed is:

1. A method for compressing concentric mosaic image data having a plurality of frames, the method comprising:

selectively dividing the plurality of frames into a plurality of anchor frames and a plurality of predicted frames;

independently encoding each of the anchor frames; and

encoding a prediction residue for each of the predicted frames, the prediction residue for each of the predicted frames being determined by referring each of the predicted frames to at least one of the anchor frames.

2. The method as recited in Claim 1, wherein independently encoding each of the anchor frames further includes:

segmenting each of the anchor frames into a plurality of anchor frame macroblocks; and

independently encoding each of the anchor frame macroblocks.

3. The method as recited in Claim 2, wherein independently encoding each of the anchor frame macroblocks further includes:

subdividing each anchor frame macroblock into a plurality of subblocks;

transforming each subblock by a discrete cosine transform (DCT); and

entropy encoding each transformed subblock using a Huffman coder.

4. The method as recited in Claim 3, wherein subdividing each anchor frame macroblock into the plurality of subblocks further includes subdividing each

anchor frame macroblock into at least one chrominance subblock and at least one luminance subblock.

5. The method as recited in Claim 3, wherein the discrete cosine transform (DCT) includes a basis-8 DCT and quantization of DCT coefficients by a quantization scale associated with the plurality of anchor frames.

6. The method as recited in Claim 1, wherein encoding the prediction residue for each of the predicted frames further includes:

segmenting the at least one anchored frame into a plurality of anchor frame macroblocks;

segmenting each of the predicted frames into a plurality of predicted frame macroblocks; and

encoding each of the predicted frame macroblocks using motion compensation.

7. The method as recited in Claim 6, wherein encoding each of the predicted frame macroblocks using motion compensation further includes:

for each predicted frame macroblock, selectively determining a significantly best match within one or more anchor frame macroblocks;

determining a reference vector for each predicted frame macroblock within each predicted frame, the reference vector indicating a position of the significantly best match within the one or more anchor frame macroblocks;

for each predicted frame macroblock, determining a prediction residue for the predicted frame macroblock by the difference between a predicted frame macroblock value and an anchor frame match value.

8. The method as recited in Claim 7, wherein encoding each of the predicted frame macroblocks using motion compensation further includes decoding each of the encoded anchor frames.

9. The method as recited in Claim 7, wherein determining the prediction residue for the predicted frame macroblock further includes:

for each predicted frame macroblock, transforming residue by a discrete cosine transform (DCT); and

entropy encoding each transformed residue using a Huffman coder.

10. The method as recited in Claim 9, wherein the discrete cosine transform (DCT) includes a basis-8 DCT and quantization of DCT coefficients by a quantization scale associated with the plurality of predicted frames.

11. The method as recited in Claim 9, wherein encoding each of the predicted frame macroblocks using motion compensation further includes using a translation-based motion model.

12. The method as recited in Claim 9, wherein encoding each of the predicted frame macroblocks using motion compensation further includes using an affine motion model.

13. The method as recited in Claim 9, wherein encoding each of the predicted frame macroblocks using motion compensation further includes using a perspective motion model.

14. The method as recited in Claim 1, further comprising outputting a bitstream having encoded anchor frame data, encoded predicted frame data, and indexing data.

15. The method as recited in Claim 1, further comprising outputting a bitstream having encoded anchor frame data associated with an anchor frame macroblock group (MBG) and corresponding indexing data.

16. The method as recited in Claim 14, further comprising outputting a bitstream that includes a thumbnail image of at least a portion of the concentric mosaic data.

17. The method as recited in Claim 14, wherein the bitstream further includes quantization scale information.

18. The method as recited in Claim 14, wherein the encoded predicted frame data includes encoded prediction residue.

19. The method as recited in Claim 14, wherein the indexing data is configured to identify each encoded anchor frame and each encoded predicted frame.

20. The method as recited in Claim 19, wherein the encoded anchor frame data is further configured to identify encoded macroblock groups (MBGs) within each encoded anchor frame.

21. The method as recited in Claim 19, wherein the encoded predicted frame data is further configured to identify encoded predicted frame macroblocks within each encoded predicted frame.

22. The method as recited in Claim 19, wherein the encoded predicted frame data is further configured to identify encoded predicted frame macroblock groups (MBGs) within each encoded predicted frame.

23. A computer-readable medium having computer-executable instructions for use in compressing concentric mosaic image data having a plurality of frames, the computer-executable instructions providing steps comprising:

selectively dividing the plurality of frames into a plurality of anchor frames and a plurality of predicted frames;

independently encoding each of the anchor frames; and

encoding a prediction residue for each of the predicted frames, the prediction residue for each of the predicted frames being determined by referring each of the predicted frames to at least one of the anchor frames.

24. The computer-readable medium as recited in Claim 23, wherein independently encoding each of the anchor frames further includes:

segmenting each of the anchor frames into a plurality of anchor frame macroblocks; and

independently encoding each of the anchor frame macroblocks.

25. The computer-readable medium as recited in Claim 24, wherein independently encoding each of the anchor frame macroblocks further includes:

subdividing each anchor frame macroblock into a plurality of subblocks;

transforming each subblock by a discrete cosine transform (DCT); and

entropy encoding each transformed subblock using a Huffman coder.

26. The computer-readable medium as recited in Claim 25, wherein subdividing each anchor frame macroblock into the plurality of subblocks further includes subdividing each anchor frame macroblock into at least one chrominance subblock and at least one luminance subblock.

27. The computer-readable medium as recited in Claim 25, wherein the discrete cosine transform (DCT) includes a basis-8 DCT and quantization of DCT coefficients by a quantization scale associated with the plurality of anchor frames.

28. The method as recited in Claim 23, wherein encoding the prediction residue for each of the predicted frames further includes:

segmenting the at least one anchored frame into a plurality of anchor frame macroblocks;

segmenting each of the predicted frames into a plurality of predicted frame macroblocks; and

encoding each of the predicted frame macroblocks using motion compensation.

29. The computer-readable medium as recited in Claim 28, wherein encoding each of the predicted frame macroblocks using motion compensation further includes:

for each predicted frame macroblock, selectively determining a significantly best match within one or more anchor frame macroblocks;

determining a reference vector for each predicted frame macroblock within each predicted frame, the reference vector indicating a position of the significantly best match within the one or more anchor frame macroblocks;

for each predicted frame macroblock, determining a prediction residue for the predicted frame macroblock by the difference between a predicted frame macroblock value and an anchor frame macroblock value.

30. The computer-readable medium as recited in Claim 29, wherein encoding each of the predicted frame macroblocks using motion compensation further includes decoding each of the encoded anchor frames.

31. The computer-readable medium as recited in Claim 29, wherein determining the prediction residue for the predicted frame macroblock further includes:

for each predicted frame macroblock, transforming residue by a discrete cosine transform (DCT); and

entropy encoding each transformed residue using a Huffman coder.

32. The computer-readable medium as recited in Claim 31, wherein the discrete cosine transform (DCT) includes a basis-8 DCT and quantization of DCT coefficients by a quantization scale associated with the plurality of predicted frames.

33. The computer-readable medium as recited in Claim 31, wherein encoding each of the predicted frame macroblocks using motion compensation further includes using a translation-based motion model.

34. The computer-readable medium as recited in Claim 23, wherein the computer-executable instructions further include the step of outputting a bitstream comprising encoded anchor frame data, encoded predicted frame data, and indexing data.

35. The computer-readable medium as recited in Claim 34, wherein the bitstream further includes quantization scale information.

36. The computer-readable medium as recited in Claim 34, wherein the encoded predicted frame data includes encoded prediction residue.

37. The computer-readable medium as recited in Claim 35, wherein the indexing data is configured to identify each encoded anchor frame and each encoded predicted frame.

38. The computer-readable medium as recited in Claim 37, wherein the encoded anchor frame data is further configured to identify encoded anchor frame macroblock groups (MBGs) within each encoded anchor frame.

39. The computer-readable medium as recited in Claim 37, wherein the encoded predicted frame data is further configured to identify encoded predicted frame macroblock groups (MBGs) within each encoded predicted frame.

40. An apparatus comprising:
memory suitable for storing concentric mosaic image data having a plurality of frames;

logic operatively coupled to the memory and configured to selectively divide the plurality of frames into a plurality of anchor frames and a plurality of predicted frames, independently encode each of the anchor frames, and encode a prediction residue for each of the predicted frames, the prediction residue for each of the predicted frames being determined by referring each of the predicted frames to at least one of the anchor frames.

41. The apparatus as recited in Claim 40, wherein the logic is further configured to segment each of the anchor frames into a plurality of anchor frame macroblocks and independently encode each of the anchor frame macroblocks.

42. The apparatus as recited in Claim 41, wherein the logic is further configured to subdivide each anchor frame macroblock into a plurality of subblocks, transform each subblock by a discrete cosine transform (DCT), and entropy encode each transformed subblock using a Huffman coder.

43. The apparatus as recited in Claim 42, wherein the logic is further configured to subdivide each anchor frame macroblock into at least one chrominance subblock and at least one luminance subblock.

44. The apparatus as recited in Claim 42, wherein the discrete cosine transform (DCT) includes a basis-8 DCT and quantization of DCT coefficients by a quantization scale associated with the plurality of anchor frames.

45. The apparatus as recited in Claim 40, wherein the logic is further configured to segment the at least one anchored frame into a plurality of anchor frame macroblocks, segment each of the predicted frames into a plurality of predicted frame macroblocks, and encode each of the predicted frame macroblocks using motion compensation.

46. The apparatus as recited in Claim 45, wherein the logic is further configured to encode each of the predicted frame macroblocks using motion

compensation by, for each predicted frame macroblock, selectively determining a significantly best match within one or more anchor frame macroblocks, determining a reference vector for each predicted frame macroblock within each predicted frame, the reference vector indicating a position of the significantly best match within the one or more anchor frame macroblocks, and for each predicted frame macroblock, determining a prediction residue for the predicted frame macroblock by the difference between a predicted frame macroblock value and an anchor frame macroblock value.

47. The apparatus as recited in Claim 46, wherein the logic is further configured to encode each of the predicted frame macroblocks using motion compensation by first decoding each of the associated encoded anchor frames.

48. The apparatus as recited in Claim 47, wherein the logic is further configured to, for each predicted frame macroblock, transform residue by a discrete cosine transform (DCT), and entropy encode each transformed residue using a Huffman coder.

49. The apparatus as recited in Claim 48, wherein the discrete cosine transform (DCT) includes a basis-8 DCT and quantization of DCT coefficients by a quantization scale associated with the plurality of predicted frames.

50. The apparatus as recited in Claim 48, wherein the logic is further configured to use a translation-based motion model to encode each of the predicted frame macroblocks using motion compensation.

51. The apparatus as recited in Claim 40, wherein the logic is further configured to output a bitstream comprising encoded anchor frame data, encoded predicted frame data, and indexing data.

52. The apparatus as recited in Claim 51, wherein the bitstream further includes quantization scale information.

53. The apparatus as recited in Claim 51, wherein the encoded predicted frame data includes encoded prediction residue.

54. The apparatus as recited in Claim 51, wherein the indexing data is configured to identify each encoded anchor frame and each encoded predicted frame.

55. The apparatus as recited in Claim 54, wherein the encoded anchor frame data is further configured to identify encoded anchor frame macroblock groups (MBGs) within each encoded anchor frame.

56. The apparatus as recited in Claim 54, wherein the encoded predicted frame data is further configured to identify encoded predicted frame macroblock groups (MBGs) within each encoded predicted frame.

57. A method for decompressing a bitstream having encoded anchor frame data, encoded predicted frame data, and indexing data associated with

compressed concentric mosaic image data having a plurality of frames, the method comprising:

accessing the index data to identify:

a unique location for each encoded anchor frame within the encoded anchor frame data and from each encoded anchor frame each encoded anchor frame macroblock group (MBG) therein, and

a unique location for each encoded predicted frame within the encoded predicted frame data and from each encoded predicted frame each encoded predicted frame macroblock group (MBG) therein; and

for each new view to be rendered:

determining which encoded anchor frame MBGs and encoded predicted frame MBGs are to be used in rendering the new view;

selectively decoding the encoded anchor frame MBG to be used in rendering the new view; and

selectively decoding the predicted frame MBG using all referenced decoded anchor frame MBGs for the predicted frame MBG.

58. The method as recited in Claim 57, wherein selectively decoding the encoded anchor frame MBG to be used in rendering the new view further includes:

for each encoded anchor frame MBG to be used in rendering the new view, determining:

if the encoded anchor frame MBG has an existing corresponding decoded anchor frame MBG, and if so, using the existing corresponding decoded anchor frame MBG in rendering the new view,

otherwise, decoding the encoded anchor frame MBG to be used in rendering the new view.

59. The method as recited in Claim 57, wherein selectively decoding the predicted frame MBG using all referenced decoded anchor frame MBGs for the predicted frame MBG further includes:

for each encoded predicted frame MBG to be used in rendering the new view, determining:

if the encoded predicted frame MBG has an existing corresponding decoded predicted frame MBG, and if so, using the existing corresponding decoded predicted frame MBG in rendering the new view, otherwise

decoding the predicted frame MBG using all referenced decoded anchor frame MBGs for the predicted frame MBG.

60. The method as recited in Claim 57, wherein each encoded predicted frame includes a prediction residue associated with at least one referenced anchor frame.

61. The method as recited in Claim 57, wherein decoding the encoded anchor frame macroblock to be used in rendering the new view further includes using an inverse discrete cosine transform (DCT).

62. The method as recited in Claim 61, wherein the inverse discrete cosine transform (DCT) includes an inverse quantization of DCT coefficients by a quantization scale associated with the plurality of predicted frames and an inverse basis-8 DCT.

63. The method as recited in Claim 57, wherein the bitstream further includes quantization scale information.

64. The method as recited in Claim 57, wherein decoding the predicted frame MBG using all referenced decoded anchor frame MBGs for the predicted frame MBG further includes:

decoding each referenced encoded anchor frame MBG for which there is no existing corresponding decoded anchor frame MBG; and

decoding the predicted frame MBG using motion compensation using a prediction residue.

65. The method as recited in Claim 57, wherein selectively decoding the encoded anchor frame MBG to be used in rendering the new view further includes storing the decoded anchor frame MBG in a first memory cache.

66. The method as recited in Claim 65, wherein selectively decoding the predicted frame MBG the referenced decoded anchor frame MBGs for the predicted frame MBG further includes storing the decoded predicted frame MBG in a second memory cache.

67. The method as recited in Claim 57, further comprising rendering the new view on at least one output device.

68. An apparatus comprising:

memory suitable for storing at least a portion of a bitstream having encoded anchor frame data, encoded predicted frame data, and indexing data associated with a compressed concentric mosaic image data having a plurality of frames; and logic operatively coupled to the memory, the logic including:

a rendering engine configured to access the index data to identify a unique location for each encoded anchor frame within the encoded anchor frame data and from each encoded anchor frame each encoded anchor frame MBG therein, and is further configured to access the index data to identify a unique location for each encoded predicted frame within the encoded predicted frame data and from each encoded predicted frame each encoded predicted frame MBG therein, and

a decoding engine that, for each new view to be rendered, determines which encoded anchor frame MBGs and encoded predicted frame MBGs are to be used in rendering the new view, selectively decodes the encoded anchor frame MBG to be used in rendering the new view, and selectively decodes the predicted frame MBG using all referenced decoded anchor frame MBGs for the predicted frame MBG.

69. The apparatus as recited in Claim 68, wherein the decoder engine is further configured to selectively decode the encoded anchor frame MBG to be used in rendering the new view by, for each encoded anchor frame MBG to be

used in rendering the new view, determining if the encoded anchor frame MBG has an existing corresponding decoded anchor frame MBG in the memory, and if so, allowing the rendering engine to use the existing corresponding decoded anchor frame MBG in rendering the new view, otherwise, decoding the encoded anchor frame MBG to be used in rendering the new view and storing the resulting decoded anchor frame MBG to the memory.

70. The apparatus as recited in Claim 68, wherein the decoder engine is further configured to selectively decode the predicted frame MBG using a decoded anchor frame MBG associated with the predicted frame MBG by, for each encoded predicted frame MBG to be used in rendering the new view, determining if the encoded predicted frame MBG has an existing corresponding decoded predicted frame MBG in the memory, and if so, allowing the rendering engine to use the existing corresponding decoded predicted frame MBG in rendering the new view, otherwise decoding the predicted frame MBG using the referenced decoded anchor frame MBG associated with the predicted frame MBG and storing the resulting decoded predicted frame MBG to the memory.

71. The apparatus as recited in Claim 68, wherein each encoded predicted frame includes a prediction residue associated with at least one referenced anchor frame.

72. The apparatus as recited in Claim 68, wherein the decoder engine is further configured to decode the encoded anchor frame macroblock to be used in rendering the new view using an inverse discrete cosine transform (DCT).

73. The apparatus as recited in Claim 72, wherein the inverse discrete cosine transform (DCT) includes an inverse quantization of DCT coefficients by a quantization scale associated with the plurality of predicted frames and basis-8 inverse DCT.

74. The apparatus as recited in Claim 68, wherein the bitstream further includes quantization scale information.

75. The apparatus as recited in Claim 68, wherein the decoder engine is further configured to decode the predicted frame MBG using the referenced decoded anchor frame MBG for the predicted frame MBG by:

decoding the referenced encoded anchor frame MBG for which there is no existing corresponding decoded anchor frame MBG and storing the resulting decoded anchor frame MBG to the memory; and

decoding the predicted frame MBG using motion compensation and a prediction residue.

76. The apparatus as recited in Claim 68, further comprising at least one output device operatively coupled to the rendering engine, and wherein the rendering engine is further configured to cause the new view to be provided to the output device.